Cell Formation and Organization in Low-Level Vision Based on Eigenpaxels

Catherine Cheung[1*] Peter McGuire[2] Gabriele D'Eleuterio[1]

[1]Space Robotics Group, University of Toronto Institute for Aerospace Studies, Toronto, ON M3H 5T6
{cathy.cheung, gabriele.deleuterio}@utoronto.ca
[2]C-CORE, St. John's, NF A1B 3X5; peter.mcguire@c-core.ca

Abstract

In this work, we propose a link between the statistical properties of natural images and cortical-cell development and organization in low-level vision. A Kohonen self-organizing neural network is used to process natural image input. The neuron weights and organized map are analyzed in terms of localized principal components, referred to as "eigenpaxels". Cortical-cell receptive-field formation, organization and function are presented in terms of these eigenpaxels.

Keywords: *low-level vision, principal component analysis, Kohonen SOM, cortical modeling, natural images.*

1 Introduction

Many efforts in the field of robotics seek to apply and emulate the biological processes that enable movement, control, and recognition. For robotic systems of the future to be truly autonomous they will require the faculty of vision, which remains to date an inadequately developed process, in part, because of the lack of thorough understanding of the biological system.

This work stems from previous efforts [17] to generate a model that replicates the logical visual strategy of the early visual cortices in mammalian species. That research addressed the similarities between principal components of image patches and cortical cell properties. The term "eigenpaxels" was employed to refer to these eigenvectors of natural image patches. In this study, we expand this model to encompass the architecture and function as well as the development of cortical cells.

The aim of this work is to further explore a correlation between eigenpaxels and the cells in low-level vision. Through this research we hope to provide insight into the principles underlying the biological vision process, and in doing so, improve the performance of vision applications in artificial systems. We therefore desire a biologically plausible but computationally elegant model. The attempt here is to explore the possibility of using fundamental engineering principles and theories to understand the brain's visual system. Eigenvectors are useful in many physical applications and in the behaviour of various natural phenomena, such as sound waves, structural vibrations, energy levels of atoms, demonstrating nature's affinity towards optimal patterns (eigenproblems can be viewed as an optimization process through the Rayleigh quotient). Here we propose another venue for eigenvectors to play a role in nature, that of the biological vision system. We present a model that seeks to explain in terms of eigenpaxels the emergence of the cortical cell receptive fields, demonstrate the properties of their known organization and architecture, and remain consistent with their function and role in the visual processing pathway. Our work in this area further supports the link established between artificial neural networks and biological vision [15], [23].

1.1 Related Work

The challenge of modeling the visual cortex and the visual processing stream has spurred a healthy number of models and theories approaching the task from various angles. Many studies concentrate on reproducing many of the cortical cell properties that Hubel and Wiesel had observed in their experiments on cats and monkeys [9], [10], [11]. Von der Malsburg put forth one of the first of these successful models which used Hebbian learning for the neural network [23], a practice which was adopted by many other researchers after him. Fukushima [6] chose to approach the task by building a neural network model with the same pattern recognition capabilities of a human in the hopes of providing a better understanding to the biological neural mechanism.

With the development of sophisticated imaging techniques [1], more detailed and complex patterns in the organization and mappings of cortical cell characteristics were discovered. Many studies focus on reproducing these cortical maps, with respect to ocular dominance

*Present Address: Institute for Aerospace Research, National Research Council, Ottawa, ON K1A 0R6

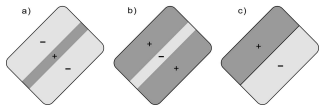

Figure 1: Simple-cell receptive-field arrangements. There are 3 arrangements of the inhibitory and excitatory regions in the RF: a) OFF-ON-OFF: oriented illuminated bar, b) ON-OFF-ON: oriented dark bar, c) OFF-ON: oriented edge. Adapted from [8].

and/or orientation selectivity trends, as in [20]. Other work seeks to address the origins of cell orientation selectivity, for example [22].

Analyzing natural images to explore their relationship with cortical-cell receptive-field structure is an approach that has been taken by several researchers. Hancock et al. [7] used a Sanger network based on Hebbian learning to extract the principal components of natural image patches in sequence through the network's multiple output units, and observed a similarity with the arrangements of cortical-cell receptive fields (RF).

The work of Olshausen and Field [19] employed a coding strategy to maximize sparseness in a set of natural images in order to obtain basis functions possessing the localized, oriented, and bandpass characteristics of simple-cell receptive fields. The images were preprocessed by filtering with a zero-phase whitening/lowpass filter. An explicit cost function was optimized during learning in the simulations.

Hyvärinen and Hoyer [12], [13] continued the sparse coding approach developed by Olshausen and Field to explain the emergence of simple-cell and complex-cell receptive-field properties as well as their columnar organization.

1.2 Early Visual Processing

Low-level vision is a subset of the form pathway, the processing stream responsible for object vision, segmenting and identifying contours, texture and colour of image data [21]. Low-level vision encompasses the retina, lateral geniculate nucleus (LGN), and the primary visual cortex. Each component of the pathway plays a role in this task. Naturally, the cells are designed and organized to accomplish these duties. Retinal cells have a centre-surround receptive field that reduces the lateral redundancy of contrast information to facilitate edge and contour detection. The cortical cells in the primary visual cortex have a receptive field assembled to perceive lines, angles, edges, and motion (Figure 1). It is these last cells, the simple and complex cells in the primary visual cortex, on which we focus our attention.

The architecture of cells in the primary visual cortex is characterized by trends in ocular dominance and orientation selectivity. These trends are generally described by Hubel and Wiesel's [8] ice-cube model and Braitenberg and Braitenberg's [2] pinwheel model. More specific features of the cortical organization are obtained from orientation selectivity and ocular dominance maps [5].

2 Low-Level Vision Model

Previous work [17] looked into the organization of Hebbian cells in a modified Földiák network when presented with natural image patches. The result was that the network's weight vectors converged to the lowest eigenvectors of the correlation matrix of pixel values formed from the image data, referred to as "eigenpaxels". The term "paxel" was used to refer to a "pack of pixels" or image patch. With raw image input, the network weights converged to the first eigenpaxel. This eigenpaxel was then removed from the input by normalizing the image, and as a result the weight vectors converged to the next two eigenpaxels.

Building on these findings, this study continues to examine the convergence of cell weight vectors with natural image input. In our model we implement a Kohonen self-organizing neural network to process the image data. The SOM is a 2-dimensional feature map which already has inherent properties that are similar to the cortex. The 2-dimensional array implies lateral connections without having to explicitly define them, and the approach of updating the winning neuron and its neighbours enables a localizing effect. The transient encountered in the modified Földiák network is avoided, and the weight update equations are in general much simpler and straight-forward.

2.1 Eigenpaxels

The principal components of natural images have an important role in this model. The eigenpaxel method was exploited in earlier work [16] on image processing. Principal component analysis is performed locally on image paxels. Image paxels of $m \times n$ pixels are randomly selected from the image set and converted to a vector $\mathbf{x}$ of length mn. A correlation matrix is formed from the image vectors

$$\mathbf{\Phi} = E[(\mathbf{x} - E(\mathbf{x}))(\mathbf{x} - E(\mathbf{x}))^T], \tag{1}$$

where $E(\mathbf{x})$ is the expected value of $\mathbf{x}$. The correlation matrix is symmetric positive-definite (except in special cases) and may be diagonalized as

$$\mathbf{\Phi} \quad = \quad \mathbf{U}\mathbf{I}^2\mathbf{U}^T, \tag{2}$$

where $\mathbf{I} = \text{diag}\{I_\alpha\}$ and $\mathbf{U} = \text{row}\{\mathbf{e}_\alpha\}$ so that I_α^2 and $\mathbf{e}_\alpha$ are the eigenvalues and eigenvectors of $\mathbf{\Phi}$ respectively.

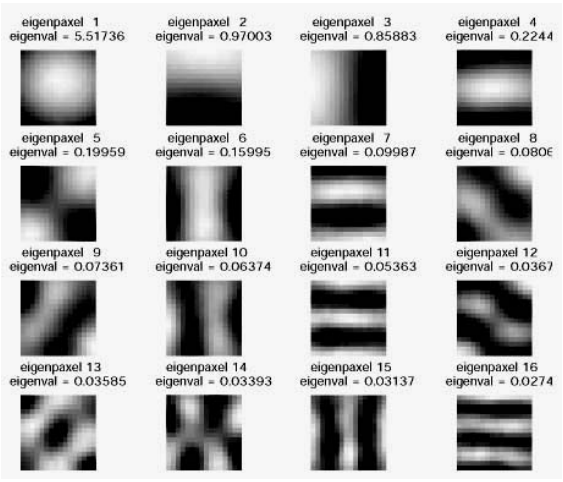

Figure 2: The first 16 eigenpaxels and their eigenvalues. The paxels are of size 16×16 pixels.

The resulting eigenvectors can be converted back into 2-dimensional paxels, or *eigenpaxels*. Figure 2 shows the first 16 eigenpaxels and their eigenvalues obtained from an image set of 1000 randomly selected paxels of size 16×16 pixels. The full images were obtained from the AT&T Laboratories Cambridge Database of Faces (http://www.uk.research.att.com/facedatabase.html).

2.2 Network Structure

The structure of the Kohonen self-organizing map (SOM) [14] that we use is a simple one-layer $M \times N$ 2-D array of cells, with nodes r_{ij}, where $i = 1, 2, .., M$ and $j = 1, 2, .., N$. Each cell or neuron has an associated weight vector $\mathbf{w}_{ij}$. Learning in this network occurs through competitive learning. For each input, $\mathbf{x}$, every cell's weights are compared to determine the closest match to the input. This comparison is carried out by minimum Euclidean distance:

$$\|\mathbf{w}_{i^*j^*} - \mathbf{x}\| \leq \|\mathbf{w}_{ij} - \mathbf{x}\|, \tag{3}$$

where i^*, j^* denotes the location of the winning cell. The weight vector of the winning cell is adjusted as well as that of the cells in the designated neighbourhood of the winner, according to the neighbourhood activity function Λ.

$$\Delta\mathbf{w}_{ij} = \beta\Lambda(|\epsilon|)(\mathbf{x} - \mathbf{w}_{ij}) \tag{4}$$

$$\Lambda(\epsilon) = \exp\left(-\frac{\epsilon^2}{2\sigma^2}\right), \tag{5}$$

where β is the learning rate, ϵ is the distance between the node and the winner, and σ is the neighbourhood size [14]. As a function of the scalar learning rate factor, $0 \leq \beta(t) \leq 1$, and the distance between the neuron and the winner, $\epsilon = r_i - r_{i^*}$, the neighbourhood function, Λ, is greatest when the node is the winner, $\epsilon = 0$, and decreases as the distance between the nodes increases. The learning rate and neighbourhood size are gradually decreased so that the weights are adjusted significantly at the beginning of training and fine-tuned as training progresses.

2.2.1 Input

We formalize the previous approach [17] approach of removing eigenpaxel components by subsequently subtracting the particular eigenpaxel projections from the input paxels $\mathbf{x}$ to obtain a filtered input paxel $\overline{\mathbf{x}}$:

$$\overline{\mathbf{x}} = \mathbf{x} - \sum_i \frac{\mathbf{x}^T\mathbf{e}_i}{\|\mathbf{e}_i\|^2}\,\mathbf{e}_i, \tag{6}$$

where $\mathbf{e}_i$ are the $m \times n$ pixel eigenpaxels, i is the number of eigenpaxel components to be subtracted, and $\|\mathbf{e}_i\|^2$ is the L_2 norm of $\mathbf{e}_i$.

The eigenpaxels are subtracted in groups of 1, 3, 6 and 10 so that eigenpaxel #1 is distinguished from eigenpaxels #2 and 3, eigenpaxels #4, 5, and 6 are considered a group, and eigenpaxels #7, 8, 9, and 10 are a set. The groupings of eigenpaxels were obtained by analyzing their eigenvalues and structure. The second and third eigenpaxels, for example, have a similar structure (Figure 2) but in orthogonal directions and therefore are treated as a pair. Eigenpaxels grouped together also have similar frequency magnitude as was shown in spectral analysis [17].

2.2.2 Implementation

The implementation details of the SOM neural network are summarized in the following steps. The series of steps listed occur for each iteration.

1. Network Input (Figure 3).

   Randomly choose an image file. The images are from a face database of 40 subjects, 10 poses each.

   From the selected image, randomly choose $m \times n$ pixel image paxel, $\mathbf{x}$.

   Calculate the eigenpaxel component(s) and subtract them from image paxel according to Equation (6).

2. Competitive Learning.

   Determine the winning node by comparing the Euclidean distance between the weights of each node, $\mathbf{w}_{ij}$, and the input paxel, $\overline{\mathbf{x}}$, according to Equation (3).

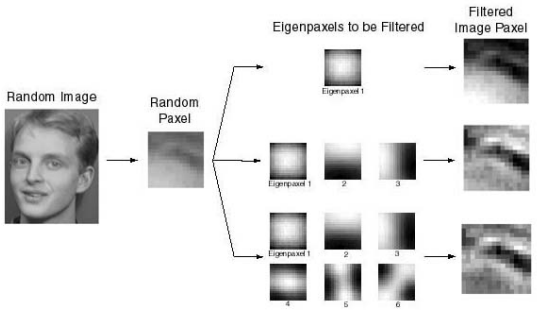

Figure 3: This figure illustrates how the filtered inputs are generated. All paxels are of size 16×16 pixels, but are not shown to scale. (1) A random image is selected from the database. (2) A random image paxel is selected from the image. (3) The eigenpaxels projections are determined and removed from the image paxel. (4) The filtered image paxel is input to the network.

Update the weights of the cells in the neighbourhood of the winning cell according to the neighbourhood activity function, Λ, using Equations (4) and (5).

Update the learning rate, β, and neighborhood size, σ,

$$\begin{aligned} \beta &= \beta_0 \exp(-3\kappa/\kappa_{max}) \\ \sigma &= \sigma_0(1 - \kappa/\kappa_{max}), \end{aligned}$$

where κ is the iteration index.

3. Weight Plots.

   Normalize the weights.

$$\mathbf{w}_{ij} = \frac{\mathbf{w}_{ij}}{\|\mathbf{w}_{ij}\|}$$

   Plot the network weights as paxels in the network array.

3 Results

Simulations were run using Matlab Release 12. Either 5000 or 10 000 iterations were run, each iteration using a different random paxel from a random face image as input. The images were stationary greyscale images obtained from the AT&T Laboratories Cambridge Database of Faces. We used 16×16 pixel-sized paxels. SOMs with fixed array sizes of 6×6, 8×8, 10×10, 12×12, and 16×16 nodes were simulated [3]. Here, however, we present only weight plots from the 6×6 and 16×16 - cell maps. All the weights were initialized to random values between 0 and 1.

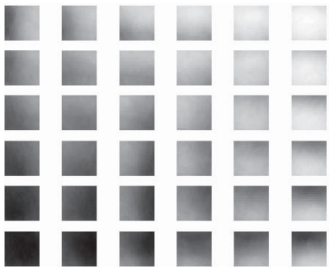

Figure 4: This figure shows the weight vectors of a 6×6 - cell SOM array with raw image input after 5000 iterations. We see a spectrum of pixel intensities in the monocontrast paxels.

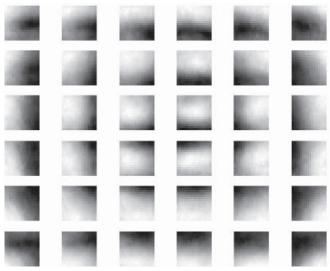

Figure 5: This figure shows the weight vectors of a 6×6 - cell SOM array with the first eigenpaxel filtered from the input after 5000 iterations. The weight paxels resemble the structure of the second and third eigenpaxels, but in all intermediate orientations.

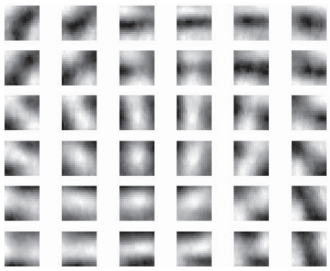

Figure 6: This figure shows the weight vectors of a 6×6 - cell SOM array with the first three eigenpaxels filtered from the input after 5000 iterations. The weight paxels are similar to eigenpaxels #4, 5, and 6 in all orientations and phases.

When the raw image paxels were input to the network, convergence was reached within 5000 iterations. Figure 4 shows the plot of the weights for a 6×6 - cell map. The first eigenpaxel is thought to represent the average pixel intensity and thus is known as the monocontrast eigenpaxel [17]. In the resulting plot, we see a palette of intensities of the monocontrast paxel, so that the distribution of the weights is along the first eigenpaxel. This result re-

produces the previous findings for raw image input [17].

In the next case, the first eigenpaxel projection is subtracted from the input image paxel. Convergence is achieved after 5000 iterations as shown in Figure 5. The weight plots resemble the structure of the second and third eigenpaxels, but in all intermediate orientations in addition to the vertical and horizontal. These weight plots show resemblance to the oriented-edge arrangement of simple cell RFs. The distribution of the oriented weights is in a pinwheel pattern around the centre.

Now when the first 3 eigenpaxel projections are removed from the image, depicted in Figure 6, the structure of the weight paxels show likeness to the next set of eigenpaxels, eigenpaxels #4, 5, and 6, but again in all orientations and phases. If we continue to subtract the next level of eigenpaxel components (#1-6), the cell maps are found to continue to distribute themselves along the next level of eigenpaxels (#7-10) [3].

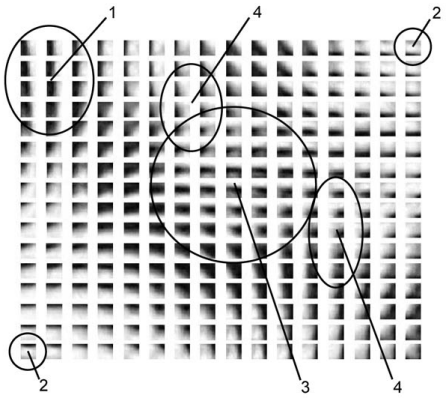

Figure 7: Orientation Selectivity Map Features on the weight plot of a 16×16 - cell SOM array with 1 eigenpaxel projection removed from the input: (1) linear zone, (2) singularities, (3) saddle point, (4) fractures.

4 Discussion

Through these simulations we verify that by using a Kohonen SOM network we are able to extract higher-order eigenpaxels from the image data when the components of lower-order eigenpaxels are removed from the input. This is not a surprising result. The ability of a neural network to perform principal component analysis has been noted and addressed by several researchers [18], [4]. As the eigenpaxels are sequentially removed from the image, the lateral competition and learning within the network allows the weights of the connections to converge to the next eigenpaxel group, distributing themselves along this eigenpaxel dimension. This allows for the development of a full spectrum of orientations which is seen in cortical cells.

4.1 Similarities to Cortical Cell Properties

4.1.1 Receptive-Field Development

First, we draw attention to the similarities between eigenpaxels (Figure 2) and cortical cell receptive fields (Figure 1). There is a striking resemblance between the first few principal eigenpaxels and the RF arrangements, as has been noted by other researchers [7], [19].

The first eigenpaxel represents the DC component of images, the average greyscale, also known as the monocontrast eigenpaxel. There is some similarity to the centre-surround arrangement, however this resemblance is attributed to the display of these eigenpaxels with heightened contrast. The second and third eigenpaxels show likeness to horizontal and vertical edges, and here we draw parallels to the oriented-line or oriented-edge structure in the cortical cell RFs. Since eigenpaxels are all mutually orthogonal and independent, the intermediate orientations are included in the space spanned by these two eigenpaxels. The fourth, fifth, and sixth eigenpaxels are reminiscent of a more distinct line structure, as opposed to an edge, although the fifth eigenpaxel is more of a checkerboard pattern with four quadrants. This set of eigenpaxels shows strong similarity to the ON-OFF-ON and OFF-ON-OFF RF arrangement. With the first five eigenpaxels of natural images, we are able to draw a remarkable likeness to the three major RF arrangements in cortical cells.

4.1.2 Organizational Properties

The results of the simulations also demonstrate properties of neural organization, the lateral spatial organization of the cells. Through lateral competition within the network, neighbouring cells show preferences for similar line orientations. This preference changes gradually from cell to cell along the map to form a configuration resembling Braitenberg and Braitenberg's pinwheel pattern [2].

In terms of the common features of orientation selectivity maps [5], there is evidence of linear zones, singularities, fractures and saddle points, as indicated in Figure 7.

As the pinwheel pattern is found in all of the simulated maps, a singularity is present at the centre. Fractures exist because orientation preferences change smoothly around the singularity and hence along a straight path there is a rapid change of orientation. Linear zones are also commonly found in these maps in regions where the orientation preference from one cell to the next does not vary greatly for linear stretches. Saddle points are not as easy to identify, but appear to develop in the larger networks near the centre of the maps at the intersection of opposing pinwheel structures.

4.1.3 Processing Features

Through the early stages of the visual pathway, the retina, LGN, and primary visual cortex, there is a sequential processing being conducted. The processing function of the retina is to adapt to local illumination levels and thereby highlight the edges in the image data [17]. This task is accomplished by passing the image through the centre-surround RFs of the ganglion cells that reduce the lateral redundancy of contrast information leaving predominantly edge information. We simulate this process by inputting raw image paxels to a Kohonen SOM network, representing the retinal cells. The processing of the network results in the weights converging to the first eigenpaxel dimension of the image, demonstrating how the cells have become sensitive to this dimension.

The processing stream continues from the retina to the primary visual cortex via the LGN. No visual processing occurs in the LGN. In the primary visual cortex, the information is passed to the cortical cells. The role of the primary visual cortex in the visual processing stream is to refine the visual information further, highlighting the line orientation characteristics of the data. Thus the cells are specific for oriented lines, and their RFs are designed to respond to these stimuli. This process is simulated by subtracting the component of the first eigenpaxel from the image and inputting this filtered image to the network, and then subtracting the components of the second and third eigenpaxels from the image for the new input. The result is of oriented lines and edges arranged topographically in the map. The cells are sensitive to lines and edges and thus are able to extract this contour information from the data, so that cells in higher cortical areas can proceed to extract more complex details such as texture and colour. Each stage then acts as a filter of the image data, removing a low-order eigenpaxel component from the image to allow the remaining features to be highlighted.

5 Concluding Remarks

While other techniques have required specific cost functions or features to be optimized, our approach is one that is less controlled. Although we detail the eigenpaxel algorithm, the active role that eigenpaxels play in the simulation is limited to filtering the input. We do begin, however, with a network that uses unfiltered raw image input. The act of filtering represents the data processing in each stage of the visual pathway. The maps that evolve then are purely the result of the network analyzing and adjusting its weights to the input image data. As similar results have been attained using a modified Földiák network based on Hebbian learning [17], the influence of the competitive learning process used here to generate these results is diminished. The network grows and tunes itself, organizing in a topographic manner and developing its individual properties through the course of information sharing between neighbouring cells. What we have shown is that a SOM will converge to the principal components of natural images, and the weights of the array will distribute themselves along these dimensions. From the organization and structure of the weight plots, we draw similarities to existing properties of cortical cell structure and organization.

As a comprehensive and functional model of the primary visual cortex, there are obviously several facets in which this model is incomplete. Architectural characteristics such as ocular dominance have not been accounted for. Cortical cell characteristics such as direction selectivity or orientation tuning bandwidths have not been addressed either. The goal of this research was to examine the extent to which eigenpaxels could be used to interpret characteristics of low-level vision. We put forth this idea in an attempt to provide insight into the neural mechanism in the brain.

Acknowledgements: This research was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

[1] G. G. Blasdel and G. Salama. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585, June 1986.

[2] V. Braitenberg and C. Braitenberg. Geometry of orientation columns in the visual cortex. *Biological Cybernetics*, 33:179–186, 1979.

[3] C. Cheung. Eigenpaxels in low-level vision: A theory of cell development, function, and organisation. Master's thesis, University of Toronto, Toronto, ON, 2002.

[4] K. I. Diamantaras and S. Y. Kung. *Principal Component Neural Networks: Theory and Applications*. John Wiley & Sons, New York, 1996.

[5] E. Erwin, K. Obermayer, and K. Schulten. Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. *Neural Computation*, 7:425–468, 1995.

[6] K. Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36:193–202, 1980.

[7] P. J. Hancock, R. J. Baddeley, and L. S. Smith. The principal components of natural images. *Network*, 3:61–70, 1992.

[8] D. H. Hubel. *Eye, Brain, and Vision*. Scientific American Library, New York, 1995.

[9] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160:106–154, 1962.

[10] D. H. Hubel and T. N. Wiesel. Receptive fields and functional architecture in two non-strate visual areas (18 and 19) of the cat. *Journal of Neurophysiology*, 28:229–289, 1965.

[11] D. H. Hubel and T. N. Wiesel. Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195:215–243, 1968.

[12] A. Hyvärinen and P. Hoyer. Emergence of complex cell properties by decomposition of natural images into independent feature subspaces. *Artificial neural networks*, 470:257–262, 1999.

[13] A. Hyvärinen and P. Hoyer. A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vision Research*, 41(18):2413–2423, 2001.

[14] T. Kohonen. *Self-Organization and Associative Memory, 2nd Ed.* Springer-Verlag, Berlin, 1987.

[15] D. Marr. Early processing of visual information. Lab Memo 340, M.I.T. Artificial Intelligence Laboratory, 1975.

[16] P. McGuire and G. M. T. D'Eleuterio. Eigenpaxels and a neural-network approach to image classification. *IEEE Transactions on Neural Networks*, 12:625–635, 2001.

[17] P. F. McGuire. *Image Classification Using Eigenpaxels*. PhD thesis, University of Toronto, Toronto, ON, 1999.

[18] E. Oja. Neural networks, principal components, and subspaces. *International Journal of Neural Systems*, 1(1):61–68, 1989.

[19] B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609, June 1996.

[20] S. J. Olson and S. Grossberg. A neural network model for the development of simple and complex cell receptive fields within cortical maps of orientation and ocular dominance. *Neural Networks*, 11:189–208, 1998.

[21] M. W. Oram and D. I. Perrett. Modeling visual recognition from neurobiological constraints. *Neural Networks*, 7:945–972, 1994.

[22] H. Sompolinsky and R. Shapley. New perspectives on the mechanisms for orientation selectivity. *Current Opinion in Neurobiology*, 7:514–422, 1997.

[23] C. von der Malsburg. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, 14:85–100, 1973.